

Tracking at CDF: algorithms and experience from Run 1 and Run II

F.D. Snider

MS-318, P.O. Box 500, Batavia, IL 60510, USA

Abstract

We describe the tracking algorithms used during Run 1 and Run II by CDF at the Fermilab Tevatron collider, covering the time from about 1992 through the present, and discuss the performance of the algorithms at high luminosity. By tracing the evolution of the detectors and algorithms, we reveal some of the successful strategies used by CDF to address the problems of tracking at high luminosities.

Key words: charged particle tracking, track reconstruction algorithms, pattern recognition algorithms

PACS: 07.05.Kf 07.77.-n 07.77.Ka 29.40.Gx

1 Introduction¹

Charged particle tracking is a critical component of almost every analysis performed at CDF. In some cases, the use of tracking is obvious and explicit, such as in the lepton identification used to isolate decays of intermediate vector bosons and J/ψ 's, or the precision vertexing of tracks in jets used to identify high-transverse momentum (p_T) b-quarks from top decays and low- p_T B mesons. Perhaps less apparent but no less important applications include the determination of the calorimeter energy scale used in various mass and jet energy measurements, the mapping of detector material using electrons from photon conversions and the triggering of the detector based in part upon the identification of a reconstructed track. The pervasive use of tracking throughout the experiment has led to a strong and continuing effort to improve the track reconstruction algorithms and to extract the best possible information from the tracking detectors.

¹ Work supported by the U.S. Department of Energy under contract No. DE-AC02-76CH03000.

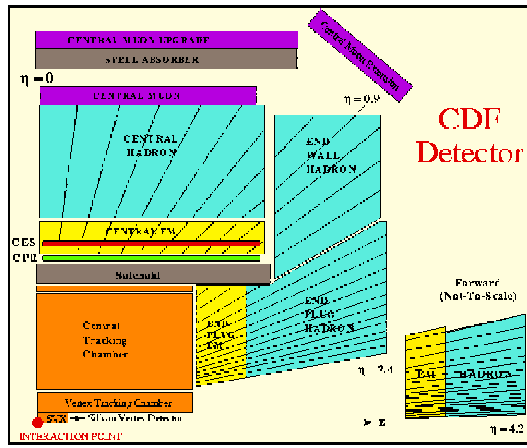


Fig. 1. Cross-sectional view of one quadrant of the Run 1 detector.

In this talk, we will review the tracking systems and the basic track reconstruction algorithms used by CDF during Run 1 and Run II at the Fermilab Tevatron collider. During this time, covering the period from 1992 through the present, improvements in the accelerator increased the instantaneous luminosity, or equivalently, the number of interactions per beam crossing by a factor of four to 15 higher than that for which the detector was designed. Consequently, tracking at high instantaneous luminosities has posed one of the greatest challenges to the detectors and algorithms. The evolution of the detector and algorithms largely reflects the history of dealing with the problems caused by high luminosity operation. By exploring the successful strategies adopted, we hope to gain some insights into more basic principles of how to perform tracking at high luminosity.

We begin by reviewing the Run 1 detector and algorithms in Sections 2 and 3. The performance of Run 1 tracking and some of the deficiencies at high-luminosity are discussed in Sect. 4. Sections 5–7 describe the same elements for the Run II detectors and algorithms. In Sect. 8, we summarize some of the more successful general strategies employed for algorithm development. Finally, we conclude in Sect. 9.

2 Run 1 tracking detectors

The Run 1 tracking system consists of three sub-detectors: the Central Tracking Chamber (CTC) [1], Vertex drift chambers (VTX) and the Silicon Vertex detector [2]. An integrated luminosity of 156 pb^{-1} was delivered to the experiment over the course of two one-year runs.




Fig. 2. Diagrams of the CTC endplate (left) showing the layout of drift cells and super-layers, and an SVX barrel (right) showing the configuration of silicon layers and detectors.

2.1 CTC

The CTC (Fig. 2) is an axial-wire drift chamber with 84 measurement layers arranged in nine alternating axial and small-angle stereo “super-layers” of drift cells. Five of the super-layers are axial and four are stereo. Each axial drift cell contains 12 sense wires, while stereo cells contain six. Sense wire planes are inclined by 45° with respect to a radial line so that the electron drift direction is azimuthal in the solenoidal magnetic field. The average single hit resolution is about $150\ \mu\text{m}$ across a maximum drift distance of about 4 cm on either side of the sense wires. Left-right ambiguities are resolved during pattern recognition.

2.2 VTX

The VTX consists of 24 modules positioned back-to-back along the beam-axis. Two planes of sense wires, segmented azimuthally into octants, measure tracks in the $r-z$ plane² at 24 different radii. Neighboring modules are offset in azimuth to provide a stereo angle, which is used to augment the stereo information in the CTC.

² The coordinate conventions used are as follows: the z -axis lies along the nominal beam axis; θ is the polar angle with respect to the z -axis; ϕ is the azimuthal angle around the z -axis; and η is the pseudo-rapidity, defined as $-\ln(\tan(\theta/2))$.

The silicon detector consists of two “barrels” (see Fig. 2) placed end-to-end around the beam pipe, each barrel consisting of 12 “wedges” with four layers of detectors located between about 2.5 cm and 10 cm from the beam. Three wafers of single-sided silicon mounted to a common support structure are used in each layer. Charge is collected by axial strips on a $60\text{ }\mu\text{m}$ pitch. The combined length of the two barrels is 50 cm, notably shorter than the luminous region. As a result, about 10% of events occur beyond the ends of the detector.

3 Run 1 tracking algorithms

Track reconstruction occurs in three phases: VTX reconstruction, during which tracks in the $r - z$ view are used to determine the z position of the primary interactions; CTC reconstruction, which itself includes several distinct algorithms and stages; and the silicon reconstruction, which is driven entirely by the results of the CTC reconstruction. Track finding in the VTX is similar to the segment finding algorithm in the CTC to be described below, and will not be discussed further.

3.1 CTC reconstruction algorithms

The CTC reconstruction proceeds in four stages: axial hit reconstruction, which occurs in two complementary algorithms; axial track merging; stereo hit reconstruction; and CTC-VTX merging. We refer to the two axial hit algorithms as Axial Segment-Linking (ASL) and Axial Hit-Linking (AHL).

3.1.1 Axial segment-linking algorithm

The ASL algorithm exploits several basic features of the axial drift cell structure. First, most tracks traverse a large fraction of the wires within a single drift cell. Second, the track trajectory within a super-layer is linear to a good approximation. Finally, the 12 sense wires in each cell is sufficiently large that robust linear segment finding is usually possible, yet small enough that the combinatoric problem is rarely difficult to solve.

Segment finding proceeds by selecting seed segments of three consecutive wires, which are then extended in both directions using a 1 mm road until all expected wires are tested. A segment search fails if more misses than hits are found. Hits possibly obscured by earlier hits on a wire are not treated

as misses. A successful search results in two segments representing the two possible drift directions.

The segment linking phase uses the interior angles between two candidate segments and the line through the center-point of each segment. For segments that are tangent to a circular trajectory, the interior angles will be equal. Linking segments using these angles corresponds to a p_T -matching criterion, which leads to the somewhat counter-intuitive result that the effectiveness of the matching criterion improves with lever-arm. The CTC algorithm links segments to a track if the angles match to within 50 mrad.

3.1.2 *Axial hit-linking algorithm*

The AHL algorithm offers a complementary approach to axial hit pattern recognition. Given the drift-cell structure of the CTC, the vast majority of tracks traverses a sense wire plan in most super-layers. Using hits that are nearest the wires in a cell, one can define a hypothetical crossing point for a candidate track. The algorithm uses this crossing point and the beam axis to define a search region. A road search is performed within this region using a road initially defined by the crossing hits, and updated as each new hit is added to the track. A candidate track is rejected if at any time during the search the number of misses exceeds the number of hits.

Both axial algorithms are performed for each event. Note that the AHL algorithm uses a strictly spatial matching criterion, which is conjugate to the p_T matching criterion of the ASL algorithm. One can easily imagine cases that are spatially dense in the detector, but when viewed in momentum space are actually sparse, and visa-versa. This type of complementarity allows the combination of the two algorithms to achieve extremely high efficiency over a broad range of operating conditions.

3.1.3 *Stereo reconstruction*

The stereo pattern recognition algorithm is performed once per reconstructed axial track. For each track, the algorithm loops over z -vertices found by the VTX. Hits in the stereo super-layers are histogrammed in $\cot(\theta)$ (see footnote 2) relative to each of these vertices. The best result, when combined with the axial track parameters, defines a 3D road. The algorithm then searches all stereo super-layers for hits within 2 cm in z of this road.

Note that unlike the two axial algorithms, the stereo pattern recognition relies upon the input of external information — it provides no stand-alone capability. We will return to this point shortly.




Fig. 3. (a) CTC efficiency as a function of p_T for low-luminosity (circles) and high-luminosity runs (squares). (b) CTC efficiency vs. pseudo-rapidity.

3.2 SVX reconstruction

The four layers of the silicon detector do not support robust stand-alone pattern recognition. Consequently, the pattern recognition algorithm simply extrapolates fully reconstructed CTC tracks into the silicon to define a search road. At each layer, the algorithm creates a new, updated trajectory for every hit within 4σ of the road by performing a progressive fit. Multiple scattering and dE/dx losses are important contributions to the road definition, and are therefore properly treated in the fit. The algorithm pursues all available branches in the road search, possibly generating many candidate tracks for each input CTC track. The candidate with the largest number of hits (with a minimum of two) and the smallest χ^2 per degree of freedom below a maximum value is defined as the reconstructed track.

4 Performance of Run 1 tracking

We characterize the performance of the tracking algorithms using the metrics of efficiency, purity and execution time, all as a function of various quantities.

The tracking efficiency for the CTC is near unity for p_T greater than about 500 MeV/c (Fig. 3a), and for $|\eta| < 1$ (Fig. 3b). Such tracks traverse all layers within the CTC. The efficiency remains above 97% in jets and is approximately constant for jets with transverse energies up to 140 GeV. A loss of about 5% is observed within the central 100 mrad of the jet core.




Fig. 4. Rate of 3D conversion electron tracks in the CTC (in arbitrary units) as a function of instantaneous luminosity.

Approximately 40% of isolated high- p_T tracks found in the CTC extrapolate to the silicon detector and traverse at least two layers. Of these tracks, $98.7\% \pm 0.1\%$ are reconstructed as silicon tracks. About half of the losses can be attributed to gaps in the detector.

We can estimate the purity of tracks by examining the fraction that point back to the vertex in a generic jet data sample. From this data, we observe that silicon hits provide a powerful confirmation of track quality. Tracks with only two silicon hits are eight times more likely to be fakes than are those with four silicon hits. The number of CTC hits on a track also predicts the fake track rate. Tracks with hits on less than half the available wires are more than twice as likely to be fakes than those with hits on almost all wires.

The usage of hits in the inner super-layer falls as a function of instantaneous luminosity from about 50% at low luminosity, to 30% at $2 \times 10^{31} \text{ cm}^{-2} \text{ s}^{-1}$. On the inner-most stereo and next axial super-layers, the usage falls by about 20% absolute over the same range. The outer-most super-layer exhibits no effect. Despite these hit losses, the axial reconstruction efficiency remains approximately constant. The efficiency for finding 3D tracks, however, falls significantly, as observed by the almost 20% drop in the relative rate of conversion electrons found as 3D tracks (Fig. 4).

Since reconstructed CTC tracks are required as input to the silicon track reconstruction, we expect the silicon tracking efficiency to fall significantly at high luminosity. We test the physics impact of this effect by measuring the relative B-tagging efficiency as a function of the number of overlaid minimum bias events. Here we observe the B-tagging efficiency falls by about 40% when five minimum bias events are added. This event density corresponds to a luminosity of about $3 \times 10^{31} \text{ cm}^{-2} \text{ s}^{-1}$ in Run 1, or about $2 \times 10^{32} \text{ cm}^{-2} \text{ s}^{-1}$ when the Tevatron is operated in the Run 2 36-bunch configuration. If we compare the effect of overlaying only the CTC data or only the SVX data, we find that the CTC contributes to upward of 80% of the efficiency loss.

The above results highlight the main deficiencies of the Run 1 tracking detectors and algorithms at high luminosity. First, the CTC contained an insuffi-




Fig. 5. Relative B-tagging efficiency in high transverse energy jet data (crosses) vs. the number of super-imposed minimum bias events. The squares (diamonds) show the result when only data from the silicon detector (CTC) is super-imposed.




Fig. 6. Cross-sectional view of one quadrant of Run II tracking detectors.

cient number of stereo layers, manifest at higher luminosities by a dramatic drop in the stereo reconstruction efficiency. Second, the silicon detector contained too few layers, lacked any 3D capability and was limited to very small radii. As a result, the silicon reconstruction was driven exclusively by the CTC, and therefore fell victim to all the ills of the CTC, such as reduced efficiency and track quality at high luminosity, and low efficiency for tracks with $|\eta| > 1$. Overall, the problems with the Run 1 tracking detectors and algorithms can be characterized as a lack of redundancy for limitations and failures in a single component.

5 Run II tracking detectors

The design of the Run II tracking attempted to address all of the major issues noted above. All tracking detectors were replaced, yet many features of the new detectors closely resembled those of their earlier counterparts. The Run II systems include the Central Outer Tracker (COT) [3], an axial wire drift chamber; and three independent silicon detectors: the Intermediate Silicon Layers (ISL) [4], the SVX II [5] and Layer 00 (L00) [6].




Fig. 7. End-view of Run II silicon detectors.

5.1 COT

Like the CTC, the COT is an open cell, axial-wire drift chamber arranged in super-layers of tilted drift cells. To accommodate higher luminosity and a shorter bunch spacing in the Tevatron, however, the maximum drift distance is 0.88 cm, about a factor of four shorter than that in the CTC. Most significantly, the number of wires in the stereo drift cells is 12, double that of stereo cells in the CTC and matching the number in the axial cells. The total number of measurement layers in the COT is 96 in eight super-layers, compared to 84 in nine super-layers in the CTC. The single hit resolution is about $140\ \mu\text{m}$, comparable to that in the CTC.

5.2 Silicon detectors

The three silicon detectors cover different radial ranges with a total of seven to eight measurement layers. The ISL, located at a maximum radius of 32 cm, contains one to two layers of double-sided silicon with a readout pitch between $110\ \mu\text{m}$ and $146\ \mu\text{m}$. SVX II covers intermediate radii between about 2.5 cm and 10.6 cm with five layers of double-sided silicon. Axial strips on the five layers have a readout pitch of $60\ \mu\text{m}$ to $65\ \mu\text{m}$, depending upon the layer. Three of the layers have 90° stereo with readout pitches between $125\ \mu\text{m}$ and $140\ \mu\text{m}$, while small-angle stereo strips on the other two have the same pitch as the corresponding axial strips. The entire detector is approximately 1 m long, thereby providing full coverage over the luminous region. The longitudinal segmentation is also more fine than that in the SVX.

L00 consists of a single layer of axial strips with a $50\ \mu\text{m}$ readout pitch located between 1.5 cm and 2.1 cm from the beam. Un-instrumented intermediate

strips allow L00 to achieve a position resolution in the range 6–7 μm .

6 Run II tracking algorithms

While the similarity of the Run 1 and Run II detectors allow considerable overlap between the algorithms used, the Run II detector improvements admit several new possibilities. Both the silicon and COT, for instance, support stand-alone 3D reconstruction. Tracks found in one detector can therefore be extended into the other. The additional silicon layers and extended coverage permit robust track-finding based upon a variety of externally generated seeds. Each of these will be described in the following.

6.1 COT reconstruction

Since the cell structure of the COT is almost identical to that of the CTC, the axial segment finding and linking algorithms draw directly from their CTC counterparts with only minor modifications to deal with the reduced maximum drift time. The similarity of the COT and CTC allowed this algorithm to be validated on Run 1 CTC data prior to the start of Run II. Most importantly, the same segment finding algorithm is now also used in the stereo super-layers, something not possible with the CTC.

While the CTC axial hit-linking algorithm could also be generalized to work in the COT, the poor scaling properties with the number of available hits and wires to search led to the development of a new, faster approach. The new technique uses axial segments found in the outer-most super-layers as seeds. For each such segment, the algorithm calculates a reference circle that is tangent to the segment and includes the beam axis. For each hit found within a large window around the reference circle, a new trajectory is calculated that goes through the middle of the segment and includes both the hit and the beam spot (see Fig. 8a). The distance from these new trajectories to the reference trajectory within some arbitrarily chosen plane within the detector is then added to a histogram (Fig. 8c). This distance effectively measures the difference in curvature between the reference circle and the circle through the hit. Real tracks lie along a single “hit circle” and therefore appear as a peak in this “constrained” curvature histogram.

A second, “unconstrained” curvature histogram is created by defining for each hit a trajectory that includes the hit and is tangent to the seed segment (Fig. 8b). The algorithm searches both histograms for track candidates, starting with the constrained histogram.

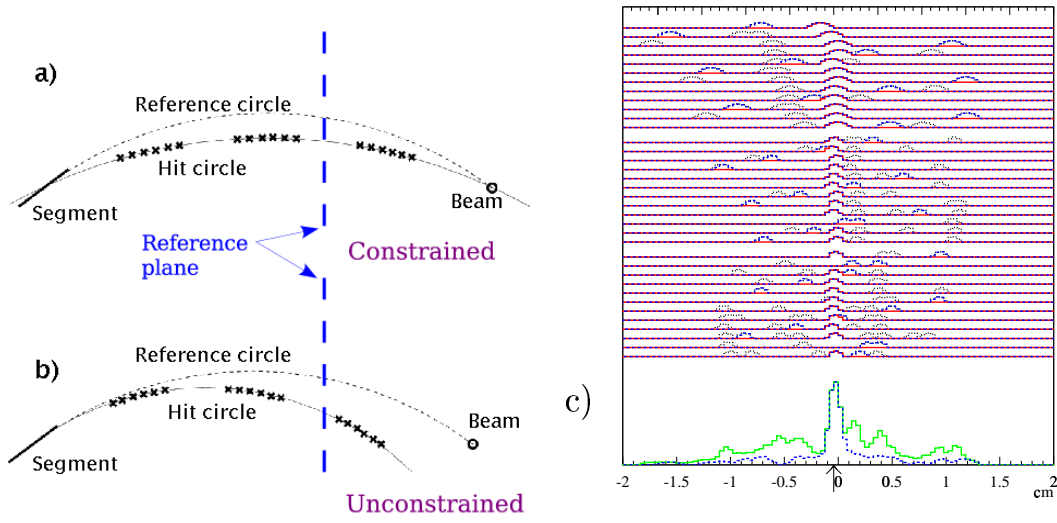


Fig. 8. The relationship between the seed segment, reference and “hit” circles, and the reference plane upon which the curvature histogram is constructed for (a) constrained and (b) unconstrained curvature histograms. (c) A curvature histogram for MC data showing (at the bottom) the histogram for all hits (solid), the correct hits (dashed), and the contribution from each axial layer (curves at the top).

The CPU time required for this axial hit-linking algorithm scales nearly linearly with hit multiplicity. It also finds high- p_T tracks with high efficiency, and is the only tracking algorithm run in the online trigger.

Stereo segment linking starts with axial tracks reconstructed in the previous steps. For each track, the algorithm performs a 3D fit using all combinations of stereo segments in the outer-most super-layers that could be consistent with the given axial track. The newly estimated 3D trajectory is then used to determine which if any of the segments in the remaining stereo layers is consistent with belonging to the track. After an exhaustive search, the combination with the best χ^2 is retained as the 3D track.

Given the resulting collection of 3D tracks, the position of the interaction vertices along the beam axis is determined using a histogramming algorithm. The vertex positions are then used to seed a stereo hit-linking phase of the type used for the CTC. Finally, the results of the segment-linking and hit-linking branches are merged and duplicates removed.

6.2 Silicon reconstruction

The standard reconstruction consists of three major phases. In the first, tracks from the COT are extrapolated into the silicon detectors, which for the purposes of the reconstruction are treated as a single unit. In this “outside-in”




Fig. 9. Seeding algorithm for silicon stand-alone reconstruction. From left to right: two candidate 3D seed hits, the $r - \phi$ and $r - z$ views of the resulting 3D seed trajectory and the nearest primary vertex (“PV”).

algorithm, axial hits are first added to the track using basically the same algorithm as that used for Run 1. The effect of material is again important and must be treated properly in defining and updating the road during the search. The candidate with the largest number of hits and best χ^2 below some maximum value is selected as a track.

Stereo hits are attached next using the same basic logic as that used for axial hits. Small-angle stereo hits are added first, followed by hits in the 90° stereo layers. After stereo reconstruction, a clean-up phase reduces potential reconstruction errors by removing hits in the inner axial layers and all the 90° stereo layers in which the chosen hit is not a significantly better choice than the second-best hit in that layer.

The second phase of silicon reconstruction consists of stand-alone pattern recognition in the silicon detectors, a new capability in Run II. It is also the most complex combinatoric problem: it is possible to have as many as 50k seeds from a few tens of tracks. Careful attention to memory and CPU management — such as caching results of repetitive and time consuming calculations, use of progressive fits, and truncating loops as early as possible — yields significant performance gains. The main goal of the algorithm is to remove candidates from further consideration at the earliest possible stage, starting with seed generation, without unduly sacrificing efficiency.

Seeds consist of the 3D trajectory through two 3D hits in the silicon (constructed by combining axial and stereo hits on the same detector) and the beam axis. The two hits must be within 9° in azimuth to be considered. Hits that belong to any other track are excluded from seed generation. The resulting helix must pass within 8 mm of a primary vertex. Seeds are then ranked according to which detector elements are used for the seed, where SVXII is ranked more highly than ISL. Given a seed, the hit search uses the same basic algorithm as for the outside-in search. Once a candidate track is accepted, all seeds that contain any hit on the track are dropped.

In the last phase of silicon reconstruction, the hits on silicon stand-alone tracks are fit in the reverse order to estimate the track parameters at the point at which the track exits the COT. This “inside-out” algorithm [7] then extrapolates the silicon track into the COT and searches for hits to attach within a road that increases from 1.4 mm to 2 mm between the inner and outer radius of the COT. The code used in this portion of the algorithm is the same as that used in both the final COT fit and a cosmic-ray finder.

6.3 *Seeded track reconstruction*

Hit finding and seed generation typically factorize in most road search algorithms. This feature allows the same hit finding code to be used with a variety of seed generation algorithms, as demonstrated with the outside-in and stand-alone silicon algorithms. Several alternative seeding algorithms, when paired with existing hit search algorithms, have proven to be highly effective pattern recognition tools and are widely used in physics analyses. The two most commonly used are the cosmic-ray finding within the COT, and calorimeter-seeded silicon tracking in the forward direction.

The cosmic-ray finder uses fully reconstructed COT tracks to generate a seed on the opposite side of the chamber. A standard hit search suitably configured then attaches COT hits to this seed. Any hits found are then fit with two assumptions, one in which the track traversed the detector from the inside out, the other that it traversed the detector from the outside-in. The same two fits are applied to the seed track. The pattern of residuals and deduced time at which the tracks pass near the beam allow cosmic rays to be readily distinguished from those created in a $p\bar{p}$ collision.

Electron tracks in the forward direction are reconstructed with high efficiency starting from the position of electromagnetic showers in the forward calorimeter. The combination of the shower and primary vertex positions and the energy of the electromagnetic shower define the helical trajectory of a hypothetical electron track to within the sign of the electron. Both seed trajectories are passed to the standard outside-in and inside-out algorithms. Typically, only a single candidate track is found for real electrons. The χ^2 of the candidate relative to the seed trajectory is used to distinguish cases in which two candidates are found.

A similar approach can be used to find hyperon tracks that traverse part of the silicon detector before decaying. In this algorithm, the decay vertex position and hyperon momentum are deduced from the decay particle tracks. The hypothetical hyperon trajectory is then used to seed the outside-in silicon algorithm. A suitable tuning of track selection parameters allows these tracks

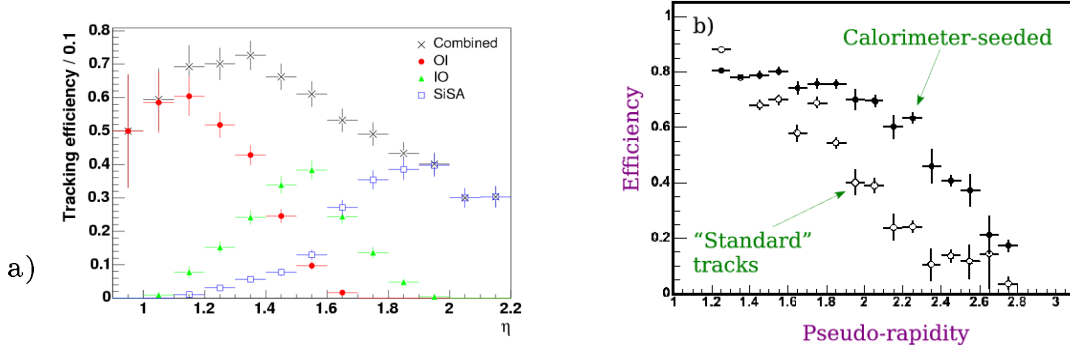


Fig. 10. Forward tracking efficiency vs. $|\eta|$ for (a) standard and (b) calorimeter seeded algorithms.

to be found in the silicon detector.

7 Performance of Run II tracking

The basic performance characteristics of the COT pattern recognition are very similar to those of the CTC, with efficiencies above 99% for isolated, high- p_T tracks in the central region. Outside-in silicon tracking has an efficiency of about 94% for attaching axial hits and 83% for stereo hits, with a fake rate below 1% for $p_T \gtrsim 3$ GeV/ c . The extended η coverage of the silicon detectors combined with the extra layers considerably improves the tracking for $|\eta| > 1$ relative to that in Run 1. Figure 10a shows the contributions to the total tracking efficiency in the forward region for each of the algorithms. The effectiveness of complementary algorithms in extending the tracking coverage is clearly evident.

Figure 10b shows the additional reach for isolated electron tracks provided by calorimeter-seeded tracking. The efficiency at $\eta = 2$ is a factor of two better than that of silicon stand-alone plus inside-out tracking. The charge mis-tag rate of these tracks is about 4%.

At high luminosity, the initial Run II COT algorithm exhibited an alarming loss of efficiency in the inner-most super-layer, which was already only 90%. An additional pass of hit association and rejection cures the problem and improves the efficiency in this super-layer to better than 98%.

The overall COT tracking efficiency for isolated, high- p_T electrons has been studied as a function of the number of reconstructed vertices in the event. We find that there is no evidence that the efficiency changes up to six vertices (Fig. 11a), which corresponds to the average number of vertices at an instan-




Fig. 11. COT tracking efficiency for (a) high- p_T electrons as a function of the number of reconstructed primary vertices, and (b) low- p_T muons as a function of the instantaneous luminosity.

taneous luminosity of about $2 \times 10^{32} \text{ cm}^{-2}\text{s}^{-1}$. Low- p_T muons show a decrease of about 1% at $10^{32} \text{ cm}^{-2}\text{s}^{-1}$ relative to that at low luminosity (Fig. 11b). The CPU time required by the tracking algorithms is approximately linear through at least $1.4 \times 10^{32} \text{ cm}^{-2}\text{s}^{-1}$. All these metrics meet the goals of the Run II detector and tracking algorithms.

8 Lessons from the CDF experience

CDF depends critically upon the reliability of the tracking detectors and algorithms. Despite large increases in the collision rate and modest increases in the maximum track density relative to Run 1, the Run II detectors and algorithms have proven successful in maintaining high efficiency and purity with relatively adequate CPU performance. This success rests less on a progression to ever more sophisticated tracking algorithms than to a fierce dedication to a few basic principles based upon the ideas of simplicity, speed, complementarity, competition, robustness and design.

All of the tracking algorithms at CDF rely upon simple road and histogram searches. The algorithms seek to reduce combinatoric complexity in stages, first solving local problems or judiciously removing data from further consideration. In this way, a complex tracking package is constructed by layering simple passes where each successive pass incorporates information from previous steps. Final decisions are deferred until the maximum amount of information is at hand.

The speed of tracking algorithms has been a primary consideration in the design and implementation. Since the execution time of the event reconstruction is dominated by tracking, imperative that they be as fast as possible. Our experience has been that simple, locally operating algorithms are the fastest.

Wherever possible, the reconstruction of no detector element or region relies upon a single algorithm. Alternative and complementary strategies have proven useful in boosting efficiency and purity without sacrificing simplicity. The tracking group also encouraged competing teams to work on the reconstruction for critical detector elements. This approach typically yielded faster code with fewer bugs than those algorithms developed without competition. In the end, the competing programs were either consolidated or used in complementary ways.

When balancing between the sometimes competing needs of efficiency and robustness, we usually opt to favor robustness. Including the wrong hits or reconstructing fake tracks is typically worse than not finding a track at all. This idea led to the development of aggressive post-reconstruction passes and final fitting programs that altered the hit content based upon a near-final track fit. Algorithms also employ constraints on the pattern recognition problem whenever possible. Beam axis or vertex constraints, hits or segments in other regions of the detector, pointing from every detector into every other, use of external detectors and stand-alone 3D capabilities are all examples of using such constraints. Tracking in dense environments is most successful when leveraging constraints.

Finally, the detector design must follow from the needs of the pattern recognition, and visa versa. None of the previous goals can be attained unless the detector and pattern recognition algorithms are designed in concert. Such considerations were important in defining the detectors to be included in the Run II tracking upgrades. The current systems build entirely upon the success of the detectors and algorithms from Run 1.

Among the many successes of the Run II detectors and algorithms, one can also find several significant short-comings. Perhaps the most evident of these is that the inner detector contains a large amount of material. This feature limits the ability to track particle at low p_T compromises the robustness of stand-alone tracking in the silicon and presents a nuisance to other analyses in the form of reduced resolution, high rate of photon conversions, etc.

A second major problem lay in the weak stereo reconstruction capabilities in the silicon. Many analyses have yet to fully integrate 3D tracking in the silicon detector. Severe quality cuts must be imposed to reduce the error rate to acceptable levels. A strong, built-in, stand-alone 3D capability would have resulted in a significantly improved tracking system.

9 Conclusions

CDF has faced a number of tracking challenges over the course of Run 1 and Run II, with detectors required to operate routinely at luminosities significantly higher than those for which the detectors were designed. Despite these challenges, the approach of using modest, incremental changes to achieve goals, building upon detectors and algorithms proven to work, and creating systems of detectors with interlocking constraints and stand-alone capabilities has produced a system that meets the physics needs of the experiment. From results in the existing data, we can now have confidence that tracking will work as needed up to instantaneous luminosities of at least $2 \times 10^{32} \text{ cm}^{-2}\text{s}^{-1}$ and beyond.

Acknowledgements

The author wishes to thank Mircea Coca, Chris Hays, Ashutosh Kotwal, Aseet Mukherjee, Rainer Wallny and Avi Yagil for useful discussions and comments in the preparation of this talk.

References

- [1] F. Bedeschi *et al.*, Nucl. Instrum. Meth. A **268**, 50 (1988).
- [2] D. Amidei *et al.*, Nucl. Instrum. Meth. A **350**, 73 (1994); S. Cihangir *et al.*, Nucl. Instrum. Meth. A **360**, 137 (1995) .
- [3] T. Affolder *et al.*, Nucl. Instrum. Meth. A **526**, 249 (2004).
- [4] A. Affolder *et al.*, Nucl. Instrum. Meth. A **453**, 84 (2000).
- [5] A. Sill, Nucl. Instrum. Meth. A **447**, 1 (2000).
- [6] C.S. Hill, Nucl. Instrum. Meth. A **530**, 1 (2004).
- [7] C. Hays *et al.*, Nucl. Instrum. Meth. A **538**, 249 (2005).